

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

WORKFLOW MANAGEMENT BASED ON AN INTEGRATED VIEW  
OF RESOURCE IDENTITY

Inventor(s):  
Stewart MacLeod  
Kim Cameron  
James Booth  
Jonathon A. Fischer

ATTORNEY'S DOCKET NO. MS1-772US

1            **TECHNICAL FIELD**

2            The following subject matter pertains to workflow management. More  
3 particularly, the described arrangements and procedures pertain to dynamically  
4 executing and managing workflow management in response to state changes in a  
5 directory.

6            **BACKGROUND**

7            All large corporations, and many small to medium sized organizations, face  
8 significant challenges in coordinating the provisioning of resources for employees  
9 and contractors. Every time a new application is installed or a new employee is  
10 hired, administrators must create multiple directory entries to match (e.g., grant or  
11 authorize access) users to the proper resources. Thus, simple events such as a new  
12 hire can turn into a laborious exercise, requiring enrollment in e-mail systems,  
13 intranet systems, remote-access systems, computers and corresponding peripheral  
14 hardware may need to be ordered, telephones provisioned, office space allocated,  
15 and the like. Similarly, if an employee leaves the company, you must judiciously  
16 remove the user from all these systems and resources.

17            In light of this, promotions, transfers, mergers, acquisitions, divestitures,  
18 marriage, divorce, name changes, and so on, can cause substantial logistical  
19 problems that can not typically be solved with a simple drag and drop of directory  
20 objects from one domain to another. Instead, users or resources must generally be  
21 deleted from one domain and then completely re-created in another domain; not to  
22 mention the logistical problems caused by moving entire groups of users and/or  
23 resources between domains. Thus, in today's fast-paced world of mergers,  
24 acquisitions, divestitures, and reorganizations, ensuring that employees have

proper access to company resources can be a full-time job for an army of information technology workers.

These resource provisioning activities traditionally occur within organizations in an uncoordinated and decentralized fashion, and often result in delays, duplication of administrative effort, and unnecessary expense. For example, the costs show up with an increasing number of calls to a help desk because directories aren't synchronized or because users can't access physical resources or applications. However, it's not just about cost and replication of effort; it's also about risk. The risks of provisioning in an uncoordinated and decentralized fashion become evident when users who are no longer with a company still exist in some of the directories and can access the company's valuable resources.

Many organizations have developed their own jury-rigged solutions to the problems of automating the provisioning process by using a workflow, which is a series of predefined actions or procedures to coordinate complex sequences of actions. Workflows have always been expensive to produce, inflexible, and difficult to maintain. For example, one problem with conventional procedures to implement workflows is that workflows are typically based on a "fire and forget" paradigm. The "fire" aspect means an individual such as a network administrator is generally required to manually identify and execute a particular workflow sequence to implement a predefined policy in response to a specific event. This is a problem because there can be any number of workflows used to implement policies within an organization. Locating a proper workflow to implement in a timely manner is not always a simple procedure and it is susceptible to human error.

1       The “forget” aspect of the fire-and-forget paradigm of conventional  
2 workflow implementation means that even though group policies are typically  
3 based on a domain or unit within an organization, no one particular entity such as  
4 a network administrator (typically has an easy way to predict which machines or  
5 resources will be affected by a workflow. This is especially true in a distributed  
6 networking environment, where peripheral devices may be added and removed  
7 from the network at any time. Additionally, because a workflow may consist of  
8 any number of long running operations that can take hours, weeks, or longer to  
9 complete, there is no easy way for an administrator to coordinate transactions  
10 and/or monitor statuses such as the success or failure of the tasks that are  
11 implemented by the workflow.

12      Moreover, once workflows have been implemented, it is typically not a  
13 simple task for an administrator to determine which particular workflow of  
14 multiple workflows was used to implement a particular policy that affected  
15 machines and/or resources.

16      For instance, consider that a particular workflow is executed to implement a  
17 policy granting one or more computers or users on a network access to specific  
18 network resources or physical facility resources to which the computers/users  
19 previously had no access. Afterward the workflow has been implemented, it is  
20 determined that one or more machines or users were erroneously granted access to  
21 those network resources or physical facilities. It may be crucial to organizational  
22 interests (e.g., security) to determine exactly which machines, users, and/or  
23 resources were affected by the workflow. However, to accomplish this it will be  
24 necessary to identify the particular workflow that resulted in the erroneous grant of

1 resources. Unfortunately, traditional workflow systems and procedures do not  
2 typically provide an easy way for an administrator to identify such information.

3 These types of problems also arise within “hire-and-fire” scenarios within  
4 an organization. For instance, after a person leaves an organization, an  
5 administrator must typically coordinate the inverse of any provisioning that  
6 occurred from when the person was hired to the time that the person left the  
7 company’s employment. The organization’s provisioning may have been  
8 implemented by any number of various workflows. Unfortunately, traditional  
9 workflow systems and procedures do not typically provide an easy way for an  
10 administrator to identify the one or more workflows that may have been used to  
11 provision the person. Thus, any modification to the person’s network and physical  
12 resource access, and any modification to associated business operations to reflect  
13 the persons new non-employee status is typically performed in an uncoordinated  
14 and decentralized fashion, possibly resulting in delays, duplication of  
15 administrative effort, unnecessary expense, and increased risk.

16 Yet another problem with traditional workflow implementations are that  
17 they are not typically “self-healing”, or self-correcting. This means that if a  
18 particular operation of a workflow fails in some manner, one or more corrective  
19 actions may not be taken in response to the failed operation. In other words, if a  
20 computer that is executing one or more aspects of a workflow malfunctions or for  
21 some other reason becomes unavailable, there is typically no mechanism to ensure  
22 that conditional aspects of a workflow that may have depended on un-  
23 implemented or failed aspects of the workflow have been satisfied. Additionally  
24 there is not mechanism to ensure that the un-implemented (or failed) aspects of the  
25 workflow are eventually implemented. This lack of a self-correcting behavior is

especially important to consider in a distributed computing environment, wherein more than one computer, and possibly dozens, hundreds, or even thousands of computers may be involved in executing any one task or operation of a workflow.

The following arrangements and procedures address these and additional problems of traditional systems and procedures to define and implement workflows.

## **SUMMARY**

The described arrangements and procedures use a directory, with its integrated view of resource identity across a distributed system to dynamically execute and manage workflow solutions responsive to changes in the directory. Specifically, a state change to an object in a directory is detected. Responsive to detecting the state change, the state change is mapped to a corresponding workflow, which includes sequences of tasks. The identified sequences of tasks are then executed to achieve a desired state in the directory. The desired state is based on the detected state change.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 shows a conventional object-oriented object class representation.

Fig. 2 shows further aspects of conventional object-oriented object class representation. Specifically, Fig. 2 shows that a class with attributes is typically represented by a rectangle divided into two regions.

Fig. 3 shows that class inheritance, or a line drawn between the subclass and the superclass conventionally represents a subclass/superclass relationship between classes.

1 Fig. 4 shows a directory schema having a flexible attribute. A flexible  
2 attribute's operational and/or data providing nature can be changed in various  
3 object instances that include the attribute without requiring directory schema  
4 modifications.

5 Fig. 5 shows an exemplary procedure to change the operational or data  
6 providing nature of multiple object instances of a base content class in a directory  
7 schema independent of modifying the directory schema.

8 Fig. 6 shows an exemplary workflow, or "provisioning" enabled directory  
9 schema to integrate workflow management with directory-based technology.

10 Fig. 7 shows further aspects of the workflow enabled directory schema.

11 Fig. 8 shows aspects of an exemplary system to manage workflow based on  
12 distributed directory technology.

13 Fig. 9 shows an exemplary procedure to manage provisioning based on an  
14 integrated view of resource identity.

15 Fig. 10 shows further aspects of an exemplary procedure to manage  
16 provisioning based on an integrated view of resource identity.

17 Fig. 11 illustrates aspects of an exemplary computing environment to  
18 manage provisioning in response to directory-based events.

19 **DETAILED DESCRIPTION**

20 The description sets forth an exemplary implementation to dynamically  
21 execute and manage a workflow using directory-based technology. The  
22 implementation incorporates elements recited in the appended claims. The  
23 implementation is described with specificity in order to meet statutory  
24 requirements. However, the description itself is not intended to limit the scope of  
25

1 this patent. Rather, the inventors have contemplated that the claimed invention  
2 might also be embodied in other ways, to include different elements or  
3 combinations of elements similar to the ones described in this document, in  
4 conjunction with other present or future technologies.

5 **Overview**

6 The described subject matter uses a directory, with its integrated view of  
7 resource identity across a distributed system to implement workflow solutions. By  
8 using the directory to coordinate the execution of tasks in a defined sequence, and  
9 to track the status of tasks within the directory until the tasks converge onto a  
10 desired directory state, the described subject matter provides a completely new  
11 way of approaching, implementing, coordinating and monitoring workflow tasks.

12 Because semantics of the described arrangements and procedures to  
13 implement a workflow are based on a workflow enabled directory schema, a  
14 schema including workflow base content classes that can be extended independent  
15 of any modification to the schema is first described. Additionally, procedures to  
16 extend the data providing and/or operational characteristics of object instances  
17 based on a directory schema independent of schema modification are described.  
18 Furthermore, a “provisioning” enabled directory schema, a system, and a  
19 procedure to integrate workflow with a directory are described.

20 **A Schema**

21 A schema is a collection of content classes and associations that abstract  
22 items, or “objects” that tangibly or intangibly exist in the real world. A content  
23 class models a set of items that have similar properties and fulfill similar purposes.  
24 A content class defines the purpose or content of an item by containing as its  
25

elements a list of properties appropriate for that purpose or content. Content classes imply a set of semantic requirements for the item. Content classes follow a hierarchical structure.

Classes can have subclasses, also referred to as specialization classes. The parent class of a subclass is referred to as a superclass or a generalization class. A class that does not have a superclass is referred to as a base class. A subclass inherits properties of its superclass. All properties and methods of a superclass apply to the subclass.

A class 10 is represented by a rectangle containing the name of the class. Fig. 1 shows an example. A class with attributes is represented by a rectangle divided into two regions as in Fig. 2, one region containing the name of the class 20 and the other region including a list of properties 22 such as what attributes are mandatory, what attributes are optional, and other properties such as what content class can be a parent of the current content class.

Class inheritance represents a subclass/superclass relationship between two or more classes. Most content classes will extend ("inherit") an existing content class. To extend a content class means that all of the properties on instances of the extended (derived) content class also exist on instances of the extending (base) content class. The act of creating an object of a particular class (or "data type") is called "instantiation" of the particular class, thereby creating an "object instance" of the class. An object instance is a collection of values, or attributes that conform to the type established by the class definition. Hereinafter, the term "object" may be used to refer to either an instance or a class.

Class inheritance can be within a namespace or across namespaces. A namespace is simply any bounded area in which standardized names can be used

1 to symbolically represent some type of information (e.g., an object in a directory  
2 or an Internet Protocol [IP] address) that can be resolved to the object itself.  
3 Inheritance is typically represented by a line drawn between a subclass and a  
4 superclass, with an arrow adjacent to the superclass indicating the superclass.  
5 Lines representing inheritance from a base class are indicated by reference  
6 numeral 30. Associations are conventionally shown as a line between two classes,  
7 as indicated by reference number 32.

8 **A Directory Schema**

9       Directory schemas are typically very carefully designed to provide content  
10 classes to meet present and future requirements of a directory. However, directory  
11 schemas are often extended to meet needs of the directory that were not  
12 foreseeable at the time that the schema was designed. For instance, just because  
13 one version of a product works with the directory schema, does not mean that  
14 other or new product versions or different products will properly function with the  
15 schema. Specifically, any variation of the type information required by a product  
16 or product versions over time generally results in the need to extend the directory  
17 schema to specifically represent each piece of interesting information that a new  
18 product or a new version of the product requires to properly operate. Because of  
19 this, third parties typically extend directory schemas to create new content classes  
20 and attributes.

21       Conventional practice, however, is to strictly control directory schema  
22 updates because modifying a directory schema requires specialized knowledge and  
23 can have complex, serious, and far-reaching consequences for customers. For  
24 example, extending directory schemas to support specific products and product  
25

1 versions means that these different products and product versions will have  
2 mutually exclusive schemas. Thus, a product that was usable with one schema  
3 may become unusable with a different schema.

4 For instance, suppose object X is an instance of class Y. Class Y has an  
5 attribute, Z. Therefore, because object X is an instance of class Y, object X can  
6 have this attribute defined on it. Assume that X does indeed have this attribute  
7 currently defined in it. Now a schema update is performed that modifies class Y  
8 by deactivating attribute Z. Note that this change makes the instance of object X  
9 invalid because X now has an attribute, Z, which X is not allowed to have  
10 according to the class definition of Y (of which object X is an instance).

11 Additionally, directory schema extensions or additions are not reversible  
12 and always add to the size of the schema. In other words, once a class or attribute  
13 has been added to the schema it cannot simply be removed from the schema once  
14 it is no longer required. Continuous schema growth due to schema extension  
15 results in a problem that is generally referred to as “schema bloat”.

16 The size of a directory schema or schema bloat becomes relevant when  
17 considering that schema changes are global to a distributed computing  
18 environment. An extended schema needs to be globally replicated to every  
19 domain server on the network. I.e., a distributed directory shares a common  
20 directory schema for the entire forest of directory trees that are organized as peers  
21 and connected by two-way transitive trust relationships between the root domains  
22 of each tree; when the directory schema is extended, the forest is extended.

23 The collection of data that must be copied across multiple servers (i.e., the  
24 unit of replication) during schema replication is the domain. A single domain may  
25 contain a tremendous number of objects (e.g., millions of objects). Thus, schema

1 extensions typically result in a substantial amount of replication traffic across the  
2 globe on multiple servers—and the larger the schema, the larger the amount of  
3 replication traffic.

4 Moreover, schema replication procedures may result in replication latencies  
5 across servers in the distributed environment, causing temporary inconsistencies  
6 between various server versions of the schema. For example, consider that a new  
7 class A is created at server X, and then an instance of this class B is created at the  
8 same server X. However, when the changes are replicated to another server Y, the  
9 object B is replicated out before the object A. When the change arrives at server  
10 Y, the replication of B fails because server Y's copy of the schema still does not  
11 contain the object A. Hence, Y does not know about the existence of A.

12 In light of these considerations, it is apparent that schema extensions  
13 typically require a substantial amount of computing resources and data bandwidth  
14 as well as coordination between network administrators to ensure that legacy  
15 applications in various domains properly operate with the updated schema.  
16 Accordingly, installing products on organizational networks that require directory  
17 schema changes can be risky, potentially politically difficult, and a time-  
18 consuming process.

19 The following arrangements and procedures address these and other  
20 problems that are associated with schema extensions by providing a workflow  
21 enabled schema with content classes that can be extended independent of schema  
22 modifications.

23 A Directory Schema with Flexible Objects and Attributes

24 Fig. 4 shows a directory schema 400 with attributes that can be extended  
25 independent of schema modifications. Specifically, the directory schema 400

1 includes a “top” or parent class 410. All other classes in the schema (e.g., the  
2 class schema class 412 and the attribute schema class 414) inherit from the top  
3 abstract class. The top abstract class includes a number of attributes (not shown)  
4 such as an X500 access control list (X500 is a well known directory protocol), any  
5 directory schema extension specific information, and other information that can be  
6 used by a directory service to instantiate the directory schema 400.

7 All directory schema 400 structural objects (other than “top”) inherit  
8 properties from the class schema class 412. Structural content classes (with the  
9 exception of the “top” content class) include only those attributes that are defined  
10 by the attribute schema class 414 or those attributes defined by content classes  
11 that have been derived from the attribute schema class 414. We now describe  
12 properties of the attribute schema content class.

13 **The Attribute Schema Content Class**

14 The attribute schema class 414 provides for a number of properties 416.  
15 Any attribute class 418 is derived from the attribute content class 414 will inherit  
16 these properties.

17 The properties 416 include:

- 18 • “*cn*”, or common-name—every object in a directory has a naming attribute  
19 from which its relative distinguished name (RDN) is formed. The naming  
20 attribute for attribute schema objects is “*cn*”, or common-name. The value  
21 assigned to “*cn*” is the value that the attribute schema object will have as its  
22 RDN.
- 23 • *LDAPDisplayName*—the name used by Lightweight Directory Access  
24 Protocol (LDAP) clients, to read and write the attribute using the LDAP

1 protocol. An attribute's `IDAPDisplayName` is unique in the schema 400  
2 container, which means it must be unique across all class schemas 412 and  
3 attribute schema 418 objects.

- 4 • `description`—a text description of the attribute.
- 5 • `adminDisplayName`—a display name of the attribute for use in  
6 administrative tools.
- 7 • `isSingleValued`--a Boolean value that is TRUE if the attribute can have only  
8 one value or FALSE if the attribute can have multiple values. If this  
9 property is not set, the attribute has a single value.
- 10 • `searchFlags`—an integer value whose least significant bits indicate whether  
11 the attribute is indexed. The bit flags in this value are: 1 = index over  
12 attribute only; 2 = index over container and attribute; 4 = add this attribute  
13 to the ambiguous name resolution (ANR) set; 8 = preserve this attribute in a  
14 tombstone object for deleted objects; 16 = copy the attribute's value when a  
15 copy of the object is created.
- 16 • `isMemberOfPartialAttributeSet`—a Boolean value that is TRUE if the  
17 attribute is replicated to the global catalog or FALSE if the attribute is not  
18 included in the global catalog.
- 19 • `systemFlags`—an integer value that contains flags that define additional  
20 properties of the attribute such as whether the attribute is constructed or  
21 non-replicated.
- 22 • `systemOnly`—a Boolean value that specifies whether only a directory  
23 service can modify the attribute.

- *objectClass*—identifies the object class of which this object is an instance, which is the class schema 412 object class for all class definitions and the attribute schema 418 object class for all attribute definitions.
- *attributeSyntax*—the object identifier of the syntax for this attribute. The combination of the *attributeSyntax* and *oMSyntax* properties determines the syntax of the attribute, that is, the type of data stored by instances of the attribute.
- *oMSyntax*--an integer that is a directory service representation of the syntax.
- *oMObjectClass*—an octet string that is specified for attributes of *oMSyntax*. For attributes with any other *oMSyntax* value, this property is not used. If no *oMObjectClass* is specified for an attribute with an *oMSyntax*, the default *oMObjectClass* is set. Usually, there is a one-to-one mapping between the *attributeSyntax* and the *oMObject* class.
- *attributeID*—the object identifier (OID) of this attribute. This value is unique among the *attributeIDs* of all attribute schema 418 objects and *governsIDs* of all class schema 412 objects.
- *schemaIDGUID*—a globally unique identifier (GUID) stored as an octet string. This GUID uniquely identifies the attribute. This GUID can be used in access control entries to control access to instances of this attribute.
- *attributeSecurityGUID*—a GUID stored as an octet string. This is an optional GUID that identifies the attribute as a member of an attribute grouping (also called a property set). This GUID is used to control access to all attributes in the property set.

- *rangeLower* and *rangeUpper*—a pair of integers that specify the lower and upper bounds of the range of values for this attribute. All values set for the attribute must be within or equal to the specified bounds. For attributes with numeric syntax the range specifies the minimum and maximum value. For attributes with string syntax the range specifies the minimum and maximum size, in characters. For attributes with binary syntax, the range specifies the number of bytes.
- *linked*—an integer that indicates that the attribute is a linked attribute. An even integer is a forward link and an odd integer is a back link.

#### An Exemplary Flexible Attribute Content Class

A flexible attribute 418 class is derived from the attribute schema content class 414. Thus, the flexible attribute inherits the properties 416. Table 1 provides an example of the values of the flexible attribute 418 class.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25

TABLE 1 – EXAMPLE OF FLEXIBLE ATTRIBUTE KEY  
PROPERTIES

Properties 416 of Flexible Attribute 418	Value
Cn	String
LDAPDisplayName	String
Description	This attribute contains XML information used by a service
AdminDisplayName	String
adminDescription	Directory ServiceInternal Use Only
isSingleValued	TRUE
SearchFlags	0x0
isMemberOfPartialAttributeSet	FALSE
Systemflags	Not Replicated
SystemOnly	FALSE
ObjectClass	Attribute Schema 414 of Fig. 4
attributeSyntax	String (e.g., XML)
OMSyntax	64
oMObjectClass	
AttributeID	TBD
schemaIDGUID	TBD

Although this example describes the flexible attribute 418 with respect to the use of XML, any other data format could be used rather than XML. For example, the attribute 418-1 could include a string in ASCII text format, or in a Hypertext Markup Language (HTML) document format, and/or the like. Significantly, the data type (e.g., “attributeSyntax”) of the flexible attribute is a text string data type. An application using an object instance that includes the flexible attribute can store, for example, an XML string on the flexible attribute property “attributeSyntax”. In this manner, the application can assign any type of information such as declarative conditions, operations, values, operational statuses, and so on, on the flexible attribute 418.

1 Accordingly, and in contrast to operational and data providing properties of  
2 conventional attributes that cannot be modified without modifying the directory  
3 schema, the operational and data providing properties of the flexible attribute 418  
4 can be changed without requiring the directory schema to be modified. To  
5 illustrate this, we now describe an object class 422 that is designed to utilize the  
6 flexible attribute class 418.

7 **An Exemplary Flexible Structural Content Class**

8 The class schema content class 412 includes object class definitions for  
9 objects 422. There can be any number of content classes 422 that are derived from  
10 class schema 412. Flexible content class 422 is derived from class schema 412  
11 and includes the flexible attribute 418. Any class that is derived from the  
12 extensible content class will inherit the flexible attribute.

13 An application using an object instance of a content class 422 can put, for  
14 example, an XML string on the flexible attribute 418. Thus, the application can  
15 assign any type of information such as data value, declarative conditions,  
16 operations, operational statuses, and/or the like, on the flexible attribute 418. This  
17 ability for an application to modify the operational and/or data providing nature of  
18 a directory object that includes the flexible attribute is accomplished without  
19 needing to modify the directory schema to create new structural object classes or  
20 attributes to include these various data and/or operational characteristics.

21 Accordingly, the described subject matter takes a new and more flexible  
22 approach to “extending” the capability of content classes that are defined in a  
23 directory schema. This is a substantial benefit over the inflexible structural  
24 content classes and attributes of conventional directory schemas.

1      **An Exemplary Procedure to Extend a Schema**

2      Fig. 5 shows an exemplary procedure 500 to change the operational or data  
3      providing nature of multiple object instances of a base content class in a directory  
4      schema independent of modifying the directory schema. At block 510, the  
5      procedure instantiates a first object instance of a flexible content class 422.

6      At block 512, the procedure assigns a first data string (e.g., XML) to a  
7      flexible attribute 418 in the first flexible object instance (block 510), the first data  
8      string defines any combination of a first operational and a data providing nature of  
9      the first object instance. Specifically, an application that has instantiated or that is  
10     using the first object instance knows of the first object instance's interface and  
11     how to unpack and use the first data string.

12     At block 514, independent of any modification to the directory schema 400,  
13     the procedure generates a second object instance of the same content class 422 that  
14     was used to create the first object instance (block 510). At block 516, the  
15     procedure puts a second data string onto the second object instance. The second  
16     data string defines a second operational and/or data providing nature of the second  
17     object instance. The first and second operational and/or data providing natures do  
18     not need to be the same. Indeed, they can be completely different in all respects  
19     other than that they are represented in a text string. The application using the  
20     second object instance knows of the second object instance's interface and how to  
21     unpack and use the second data string.

22     For instance, consider that an application can assign the flexible attribute in  
23     the first instantiated object to have any combination of one or more data types  
24     (e.g., integer, real, string, floating, character, and so on), or operational properties  
25     (e.g., an operation can be defined to do just about anything imaginable such as to

send an e-mail message, to report statistics, to manage a rocket launch, and so on). Whereas the flexible attribute in the second instance of the object can be assigned completely different properties that are independent of any characteristics of the data types or operations that correspond to the flexible attribute of the first instance of the object.

In yet another example, consider the following XML string shown in Table 2.

---

**TABLE 2**  
**EXAMPLE OF A FIRST STRING TO BE APPLIED TO A FIRST INSTANCE OF THE FLEXIBLE OBJECT**

```
<data>
    <dataType>integer</dataType>
        <value>2</value>
        <name>integerValue</name>
    <DataType>Real</DataType>
        <value>512.6</value>
        <name>realValue</name>
    <dataType>integer</datatype>
        <name>result</name>
</data>
<operation>
    <integerVal + abs(realValue) = result>
</operation>
```

---

An application can assign the string of Table 2 to a flexible attribute 418 of type string in a first instance of a flexible object 422. In this case, the string of Table 2 provides both data and operational properties to the flexible attribute. Specifically: (a) an integer variable “integerValue” is defined with a value of two (2); (b) a real variable “realValue” is defined with a value of 512.6; and (c) an addition operation that adds integerValue to the absolute (“abs”) value of realValue is defined. Thus, the XML string of Table 1 provides specific data and operations to the first instance of the object.

Now consider the following XML string of Table 3.

**TABLE 3**  
**EXAMPLE OF A SECOND STRING TO BE APPLIED TO A SECOND**  
**INSTANCE OF THE FLEXIBLE OBJECT**

<application>www.somedestination.org/applicationname.exe</application>

Independent of any modification to the directory schema, the application can assign the string of Table 3 to a flexible attribute 418 of type string in a second instance of a flexible object 422. In this case, the string of Table provides data. Specifically, the string identifies a Universal Resource Location (URL) of a computer program application.

In contrast to conventional schemas (wherein once an attribute is assigned a particular data type, only data of that predetermined data type can be represented by that attribute), a flexible attribute 418 can take on multiple values (e.g., integers, real numbers, operations, and so on) independent of the attributes 418 actual data type. Specifically, as the previous examples show, the described arrangements and procedures accomplish this independent of modifications to the directory schema to create corresponding content classes.

This multi-valued aspect of a single attribute of multiple object instances of the same base content class in a directory schema, allows a directory schema to be “versioning aware”. Specifically, this is because application providers can upgrade and provide new products that utilize flexible attributes 418 without extending the directory schema. This allows third parties to provide products and product upgrades without extending directory schemas to take into consideration the specific needs of the products and product upgrades. Accordingly, a directory

1 that is based on a directory schema 400 comprising object classes 422 with  
2 flexible attributes 418 is a “versioning aware” directory.

3 For example, consider the first XML string or document “`<a> Data </a>`”.  
4 A first version of a product understands and extracts this first string. A third party  
5 or user can simply extend the first document to support additional product versions  
6 or a new product by appending new data to the original data. For example, the  
7 following information: “`<b>Data2</b>`” can be appended to the first document to  
8 obtain the following: “`<a> Data </a><b>Data2</b>`”. In this case, the original  
9 data format of the first string is maintained and the first product versions (e.g.,  
10 legacy applications) are able to continue operations using the original data format.  
11 New applications or product upgrades that are aware of new data (e.g., “`<b>`”) can  
12 obtain the new data from the document.

13 Accordingly, while extending the data characteristics and/or operational  
14 functionalities of various object instances of a same directory schema base content  
15 class, the described arrangements and procedures completely avoid schema bloat  
16 as well as the complex and serious consequences of procedures to extend and  
17 replicate a directory schema because the directory schema is not modified.

18 **An Exemplary Workflow Enabled Directory Schema**

19 Fig. 6 shows an exemplary workflow enabled directory schema 600 with  
20 flexible objects 612 and attributes 620 that integrate workflow management with  
21 directory-based technology. Specifically, the provisioning enabled directory  
22 schema 600 includes a “top” content class 410. All other content classes in the  
23 schema such as class schema base content class 412 and attribute schema base  
24 content class 414 inherit from the top class. Top includes a number of attributes  
25

1 (not shown) such as an X500 access control list (X500 is a well known directory  
2 protocol), any directory schema extension specific information, and other  
3 information used by a directory service to instantiate objects based on the  
4 directory schema 600.

5 All directory schema 600 structural objects (other than “top”) inherit  
6 properties from class schema class 412. For instance, provisioning structural  
7 content classes 612 derive from flexible base content class 422 and class schema  
8 412. These novel classes 612 are used to execute and manage provisioning in a  
9 distributed computing environment based on the dynamic detection of state  
10 changes in directory objects. These provisioning content classes are described in  
11 greater detail below in reference to Fig. 7.

12 Structural content classes can only include attributes (with the exception of  
13 “top” 610) that are defined by attribute schema class 616 or by classes derived  
14 from the attribute schema class. The attribute schema base content class 414  
15 includes a number of properties 416. This attribute base content class 414 and  
16 these properties 416 have been described in detail above. Any attributed classes  
17 620 derived from the attribute schema base content class 414 inherit these  
18 properties 416. For instance, provisioning attributes 620 derive from flexible  
19 attribute base content class 418, which in turn derives from the base attribute  
20 schema content class 414.

21 **Provisioning Attribute Classes 620**

22 The provisioning attribute content classes 620 include the provisioning  
23 status attribute 624, the provisioning XML status attribute 626, the provisioning  
24 configuration attribute 628, the provisioning configuration private attribute 630,  
25

the provisioning status binary attribute 632, the provisioning link attribute 634, and the provisioning back link attribute 636. Provisioning attribute 620 content class characteristics allow declarative conditions, operations, operational statuses, and so on, to be stored on object instances at any time to obtain various different behaviors and data types according to an application's needs. Thus, individuals can store various data type information onto a provisioning object instances (i.e., object instances provisioning classes 614) which include a flexible attribute. Although we show each of these attributes 620 as being derived from the flexible attribute, not all need to be. For instance, one or more of these provisioning attribute content classes may not be derived from the flexible attribute 418.

The Provisioning Status Attribute 624

An instance of the provisioning status attribute 624 stores the status information that corresponds to a workflow object 510 of Fig. 5. This storage is optional because there may be a policy about retaining this information (e.g., the information may not be retained because of a data storage issue). This status information is stored in an internal structure that contains the current execution status of the provisioning rules associated with this entry. The provisioning status attribute 624 is not indexed, replicated or included in a global catalog, and is only stored on the domain controller that the directory service instance is running on.

For example, an object instance that includes the provisioning status attribute is an XML document that includes the status of a provisioning policy for an entry in a directory. Although this example and some of the following examples describe the use of XML, any other document format could be used as well as than XML. For example, an instance object that includes the provisioning

1 status attribute could store a document in ASCII text format, the Hypertext  
2 Markup Language (HTML) document format, and/or the like, on the attribute.

3 The Provisioning Status XML Attribute 626

4 The Provisioning-Status XML attribute 626 is one or more documents that  
5 store workflow based status information. This information is represented in a text  
6 string such as an XML representation of the status of a workflow, and will be  
7 stored on structural object instances of class provisioning status 512 of Fig. 5.  
8 This attribute is not indexed, replicated or included in the global catalog. It is only  
9 stored on the domain controller that the directory service instance is running on.

10 The Provisioning Configuration Attribute 628

11 An instance of the provisioning configuration attribute 628 content class  
12 stores one or more provisioning policies, or rules in a document data format such  
13 as in an XML data format. An instance of the provisioning configuration attribute  
14 628 stores the provisioning policies as a document such as an XML document.  
15 This attribute is not indexed, replicated or included in the global catalog. It is only  
16 stored on the domain controller that the MMS service instance is running on.

17 The Provisioning Configuration Private Attribute 630

18 An instance of the provisioning configuration private attribute 630 content  
19 class provides an option to store private provisioning status information. In this  
20 implementation, the information is encrypted and only accessible by  
21 administrators by default. This attribute is not indexed, replicated or included in  
22 the global catalog, and is only stored on the domain controller that the directory  
23 service instance is running on.

1                   The Provisioning Status Binary Attribute 632

2                   The provisioning status binary attribute 632 is an octet string that stores a  
3 summary of status of executing workflows in the directory. This information is  
4 stored in an internal structure and consists of the object's globally unique ID  
5 (GUID) of the workflow, the GUID of the workflow step, and the status of the  
6 associated task. Whenever a new workflow object 510 of Fig. 5 is assigned to a  
7 provisioning status object 512, the information is added to the attribute and  
8 modified according responsive to changes in status. This attribute is not indexed,  
9 replicated or included in the global catalog.

10                  The Provisioning Link Attribute 634

11                  An instance of the provisioning link attribute 634 resides on a  
12 metadirectory object and indicates or references instances of the provisioning  
13 status structural content class 512 of Fig. 5, which is described in greater detail  
14 below. The provisioning link is a single-valued domain name (DN) that points to  
15 the associated entry in a directory partition that is being used to track the status of  
16 workflow (i.e., that is being used to store instances of the provisioning objects 510  
17 through 520). An instance of the provisioning service structural class 518 of FIG.  
18 5 writes this value to each user, group or organizational unit (OU) under the  
19 control of directory enabled workflow. In this manner, the attribute provides a  
20 backlink to automatically maintain inter-object relationships in the directory.

21                  The provisioning link attribute 634 is not indexed, replicated or included in  
22 the global catalog. It is only stored on the domain controller that the directory  
23 service instance is running on.

1                   The Provisioning Backlink Attribute 636

2                   The provisioning backlink attribute 636 resides in the status object and  
3 provides a backlink that helps track an object that a provisioning agent 618 of has  
4 performed work on (e.g., executing one or more operations that corresponds to of  
5 a workflow object). The provisioning agent writes this value to each user, group  
6 or organizational unit (OU) under the control of directory enabled workflow. This  
7 attribute is not indexed, replicated or included in the global catalog. It is stored on  
8 the domain controller that the directory service instance is running on.

9                   The Class Schema Structural Object Class

10                  Fig. 7 is a block diagram that shows further aspects of the provisioning  
11 structural classes 614 that are derived from class schema 612 of Fig. 6. The  
12 provisioning structural classes include: the workflow class 710, the provisioning  
13 status class 712, the provisioning service history class 714, the service connection  
14 point class 716, the provisioning service class 718, and the event association class  
15 720.

16                  Instance objects that correspond to these provisioning classes 614 can be  
17 stored in a domain-naming context (DNC) or in an application-naming context  
18 (ANC). However, in this implementation the provisioning instance objects are  
19 stored in a directory partition such as the ANC that is not necessarily, but rather  
20 optionally replicated to other domain controllers in a forest. Instances of  
21 workflow objects 710 and provisioning status objects 712 are stored in a connector  
22 space in a directory partition such as the ANC.

23                  We now describe exemplary aspects and attributes of these classes.  
24  
25

1                   The Workflow Structural Class 710

2         The workflow structural class 710 is used to store data format  
3         representations (e.g., XML representations) of workflows in a connector space  
4         such as the described ANC space. An instance of the workflow class 710  
5         indicates a sequence of actions that correspond to a workflow. A workflow's  
6         sequences of actions are implemented and persisted within a system (i.e., a system  
7         600 of Fig. 6 by a corresponding workflow process 622 that is based on class 710).

8         Table 4 illustrates an example of a workflow. In this implementation, a  
9         workflow is represented using XML.

---

10                   **TABLE 4**  
11                   **EXAMPLE OF A WORKFLOW REPRESENTATION WITH A**  
12                   **PRECEDENCE CONSTRAINT**

13                   <?xml version="1.0"?>  
14                   <Workflow>  
15                    <WorkflowStep name="NotifyStart" ID="1">  
16                    <task type="process">...</task>  
17                    </WorkflowStep>  
18                    <WorkflowStep name="IWorkflow1" ID="2">  
19                    <task type="COM">...</task>  
20                    <PrecedenceConstraints>  
21                    <constraint ID="1" type="success"/>  
22                    </PrecedenceConstraints>  
23                    </WorkflowStep>  
24                   </Workflow>

---

25         The elements that comprise the workflow of Table 4 include a series of  
26         steps ("WorkflowStep") that identify two tasks, the "NotifyStart" task and the  
27         "IWorkflow1" task. The "NotifyStart" task has a "task type" that is a "process".  
28         The "IWorkflow1" task has a task type" that is a COM-based task. The workflow  
29         indicates that the "IWorkflow1" task has a precedence constraint

1 (“PrecedenceConstraints”) that requires the “NotifyStart” task to return with a  
2 status of “success” before the workflow will invoke the “IWorkflow1” task. This  
3 provides for the representation of any arbitrary sequence of steps in a workflow as  
4 a directed a-cyclic graph—any complex series of task based relationships, task  
5 status enumerations, and the like, can be represented. Such relationships include,  
6 for example, a task finish-start relationship, a parallel task execution relationship,  
7 simple and complex precedence constraint relationships, task priority  
8 relationships, and the like. Precedence or scheduling constraints are based on  
9 provisioning task success, failure, and completion determinations.

10 A workflow schema (i.e., the workflow schema 634 of Fig. 6) enforces the  
11 semantics of a workflow indicated by an instance of the workflow class 710.

12 Traditional techniques to implement a workflow typically require a  
13 workflow to schedule a “constrained” and/or “prioritized” task by *looking forward*  
14 to determine whether such constraints/priorities have been met. This means that a  
15 traditional workflow typically looks forward to analyze task constraint/priority  
16 relationships to determine if tasks should be executed. This is a problem because  
17 the complexity of a workflow increases at a rate that is substantially proportional  
18 to a square-root of the number of tasks that the workflow is coordinating.

19 In contrast to a traditional workflow implementations that typically require  
20 a workflow task to manage a constrained and/or prioritized task(s) by *looking*  
21 *forward* to determine the status of a condition, a workflow object 710 represents  
22 semantics of task precedence constraints and/or priorities by *looking backwards*,  
23 rather than looking forward. A task looks backward to determine if any  
24 constraints, conditions, priorities, and so on, have been satisfied prior to  
25 implementing its portion of the workflow. This substantially simplifies the

1 complexity of representing, determining, and implementing task  
2 precedence/priority relationships. For example, task D looks back at task C to  
3 determine if C completed successfully before task D executes.

4 Moreover, a workflow object 710 provides for the representation of a  
5 substantial range of workflow task types and information. For example, a task can  
6 be a process, a Component Object Model (COM) object, identified with a  
7 Universal Resource Locator (URL), a Simple Object Access Protocol (SOAP)  
8 object, a message queue, an Microsoft Transaction Server ® (MTS) transaction,  
9 and so on. (COM is a well-known binary code developed by Microsoft  
10 Corporation of Redmond, WA. (SOAP is a well-known protocol that provides a  
11 way for applications to communicate with each other over the Internet,  
12 independent of platform).

13 The Provisioning Status Structural Class 712

14 A provisioning status structural class 712 instance tracks statuses of long-  
15 running workflows that may be executing on any number of different computers.  
16 This tracking is done from a centralized location in a distributed directory. To  
17 track workflow statuses, the object 712 may contain instances of the described  
18 provisioning link attribute 634 of Fig. 6, the provisioning status attribute 624,  
19 and/or the provisioning status binary attribute 632. These attributes allow a  
20 provisioning status structural object to store information such as XML and binary  
21 representations of workflow statuses.

22 An instance of the provisioning status object 712 is generated and used by  
23 an instance of the provisioning service/agent class 718 (see also, the provisioning  
24 service 618 of Fig. 6) to monitor a workflow's status, and to reflect the current and  
25

1 past conditions within a directory that are attributed to the workflow (e.g., as any  
2 directory resources affected by the workflow). The properties of this class indicate  
3 the sequences of actions that makeup a workflow along with status information  
4 that corresponds to the sequences. Detailed examples of how the provisioning  
5 service uses one or more provisioning status objects to track workflow are  
6 provided below in reference to Figs. 6-8.

7                   The Provisioning Service History Structural Class 714

8                   An instance object of the provisioning service history structural class 714 is  
9 used by an instance object of the provisioning service/agent class 718 (see also,  
10 the provisioning service 818 of Fig. 8) to store the history of provisioning policies  
11 for an organization. This class does not contain any security principal attributes. .

12                  In one implementation, an object instance of the provisioning service  
13 history class 714 is stored in a container of a specified directory partition in the  
14 directory such as a container named “Provisioning Agent History”.

15                  Such organizational provisioning history status can be used by a data  
16 transformation service to record the history of package execution and data  
17 transformations (i.e., data lineage) for in the status information. .

18                   The Service Connection Point Structural Class 716

19                  An instance of the service connection point structural class 716 represents  
20 an instance of a service object running on one or more computers on a network.  
21 (A service object is a program, routine, or process that performs a specific system  
22 function to support other programs). The class 716 provides for a service object to  
23 explicitly publish itself in a directory such as Active Directory ®, thereby  
24 facilitating service-centric administration and usage.

1                   The Provisioning Service Structural Class 718

2                   The provisioning service/agent class 718 inherits from the connection point  
3                   class 716. An instance object of the service/agent class 718 is a service object that  
4                   executes on one or more computers on a network. Specifically, the provisioning  
5                   service is a program module that coordinates and tracks workflow/provisioning  
6                   policies in a distributed computing environment using directory-base technology.

7                   To accomplish this, the service 718 monitors directory-based events as well  
8                   as other events across the network and maps the detected events to defined  
9                   workflows. An enumerated list of directory-based events is substantial because it  
10                  includes all events that are propagated by a modification not only to the directory  
11                  schema, but also to any objects or attributes defined by the directory schema. For  
12                  example, directory-based events include creating a new directory object, deleting a  
13                  directory object, renaming a directory object, synchronizing a directory object  
14                  with an offline directory object, replicating a directory schema, and/or the like.  
15                  (Detailed examples of how the provisioning service uses one or more provisioning  
16                  status objects to track workflow are provided below in reference to Figs. 6-8).

17                  The provisioning service also keeps track of all the attributes and locations  
18                  of shared provisioning objects (i.e., instances of provisioning content classes  
19                  710 through 720) across the distributed computing environment through its  
20                  directory-based integrated view of resource identity.

21                  The service/agent 718 explicitly publishes itself in a directory (via a  
22                  corresponding connection point object 716) such as Active Directory® to  
23                  facilitate service-centric administration of workflows in an organization.

24                  A user that wants to communicate with an object instance of the  
25                  provisioning service 718 may interrogate the “explicitly published” provisioning

1 service using namespace syntax (e.g., “//workstation/provisioning agent”). Or, the  
2 provisioning service 718 may provides an abstraction layer (e.g., a dynamic-link  
3 library (DLL)) to hide the particular service location details within the distributed  
4 computing environment from client applications. The abstraction could query the  
5 directory service for a connection point object 716 representing the provisioning  
6 service/agent and use the binding information from that object to connect the  
7 client application to the provisioning service.

8 The Event Association Structural Class 720

9 The Event-Association structural class 720 is used to map declarative  
10 workflow conditions to particular instances of workflow objects 710. These  
11 provisioning-rule objects are stored in a container named after a provisioning  
12 service in the domain partition.

13 An instance object of the event association class 720 indicates an XML  
14 representation of directory events and maps them to already defined (although not  
15 necessarily instantiated) workflow (i.e., see workflow objects 710 of Fig. 7) based  
16 on a set of categorization rules that indicate declarative conditions (the  
17 categorization rules map an event to a particular workflow). A second property of  
18 the event association attribute indicates a distinguished name (DN) of the  
19 particular workflow that is associated to the detected event. The DN also  
20 identifies the workflow object’s location in a directory tree of a distributed  
21 directory.

22 An event association object’s 720 declarative conditions are based on a type  
23 of event (e.g., an add event, a delete event, a rename event, and the like), a scope  
24 of the event (what portion of the directory tree the event occurred in), any filtering  
25

1 criteria (such as the structural content classes 710 through 720 of Fig. 7 or  
2 attribute values that are pertinent to the event), and the like.

3 To illustrate such event categorization rules, consider the following  
4 example, wherein an XML script is used to determine if a detected event (a  
5 directory-based event or otherwise) maps to a defined workflow (i.e., an  
6 instantiated workflow object 712) is shown in Table 5.

7

---

8 **TABLE 5**  
**EXAMPLE OF EVENT EVALUATION TO DETERMINE A**  
**DEFINED WORKFLOW**

9 <?xml version="1.0"?>  
10 <!-- Fire if: it is an ADD event and the objectClass is PERSON  
11 AND the title attribute is EITHER PROGRAM MANAGER OR  
12 PRODUCT MANAGER-->  
13 <eventAssociationDoc>  
14 <expression>  
15 <logicalop optype="and">  
16 <event>add</event>  
17 <objectClass>person</objectClass>  
18 <logicalop optype="or">  
19 <attrConstraint name="title" comparator="equals">Program Manager  
20 </attrConstraint>  
21 <attrConstraint name="title" comparator="equals">Product Manager  
22 </attrConstraint>  
23 </logicalop>  
24 </logicalop>  
25 </expression>  
</eventAssociationDoc>

---

26 The script's declarative conditions first identify the detected event's type to  
27 determine if it is an "add" event. (The scope of the event (the portion of the  
28 directory tree the event occurred in and which directory objects are/were effected  
29 by the event) is provided by the directory service 816 of Fig. 8). If the event is an  
30 "add" event, the script determines whether an object associated with the add event  
31

1 has a corresponding content class of “person”. If so, the script determine if the  
2 object having a class of person has a corresponding “title” attribute of either  
3 “program manager” or “product manager”.

4 In this example, if these declarative conditions are met, then a  
5 corresponding workflow (indicated by the workflow indication attribute) is  
6 executed. (The workflow structural content class 710 and workflow semantics are  
7 described above (see, also the workflow schema 828 and workflow processes 826  
8 of Fig. 8).

9 The semantics of an event association object’s declarative conditions  
10 (e.g., the conditions shown in Table 5) are enforced by a schema such as an event  
11 association schema 832 of Fig. 8, which is based on the World Wide Web  
12 Consortium (“W3C”) schema standard.

13 In yet another example of an event association object’s 720 corresponding  
14 event categorization rules, consider the XML script shown in Table 6, which  
15 determines if a detected “directory modification” event maps to a defined  
16 workflow object 710.

**TABLE 6**  
**EXAMPLE OF EVENT EVALUATION TO DETERMINE IF A  
 PARTICULAR WORKFLOW IS TO BE IMPLEMENTED**

```
<?xml version="1.0"?>
<!-- Fire if:
it is a MODIFY event and
the employeeType attribute is REGULAR FULL TIME
AND the startDate attribute is less than July 6, 1998
-->
<eventAssociationDoc>
  <expression>
    <logicalop optype="and">
      <event>modify</event>
      <attrConstraint name="employeeType"
        comparator="equals">Regular Full Time
      </attrConstraint>
      <attrConstraint name="startDate"
        comparator="lt">1998-07-01</attrConstraint>
    </logicalop>
  </expression>
</eventAssociationDoc>
```

The declarative conditions shown in Table 6 first identify the detected event’s type to determine if it is a “modify” event. If so, the script determines whether an object that corresponds to the modify event has an employee type attribute (“employeeType”) indicating a regular, full time employee status (“Regular Full Time”). If so, the script determines if the object has an attribute that indicates a hire, or start date (“startDate”) before a predetermined starting date (i.e., “1997-08-01”). In this example, if these declarative conditions are met, then a corresponding workflow identified by an associated workflow attribute (not shown) of the event association object is executed.

In this implementation, an event association's 720 declarative conditions are represented in XML and stored in a directory partition that is optionally

1 replicated to other domain controllers. In yet another implementation, event  
2 association objects are represented as special structures in memory such as random  
3 access memory (RAM) for system performance reasons.

4 **An Exemplary System**

5 Fig. 8 shows an exemplary system 800 to integrate workflow management  
6 with distributed directory technology. The system provides a logically centralized,  
7 but physically distributed view of workflow that is based on a distributed directory  
8 infrastructure of machines and resources. In response to detected directory-based  
9 events such as changes in the directory infrastructure, the system implements,  
10 coordinates, and/or monitors a corresponding workflow across many different  
11 machines.

12 For example, in response to detected changes in a directory infrastructure,  
13 the system 800 automatically executes workflow tasks based on any precedence  
14 constraints and/or priorities between workflow tasks, provides workflow task  
15 status, identifies resources affected by a workflow, and/or provides for automatic  
16 implementation of corrective action in response to a workflow task failure. Thus,  
17 the system dramatically simplifies procedures to provision machines and resources  
18 such as user accounts, software configuration and updates, resource allocation, and  
19 the like.

20 Additionally, the system 800 executes a workflow until the convergence of  
21 a desired state of a number of objects in the directory is achieved (e.g., with  
22 respect to a group of directory objects representing tangible and intangible  
23 objects). For instance, each task in the workflow has one or more operations to

1 implement. Upon failure of any operation, the task may retry the operation until  
2 the desired state is achieved.

3 To accomplish this, the system 800 includes a provisioning computer 802,  
4 which is turn is a domain controller for managing user access to distributed  
5 directory resources. The provisioning server streamlines internal organizational  
6 processes by distributing a workflow between other servers 806 and expediting  
7 computing processes by harnessing the power of the other servers. The  
8 provisioning server 802 is operatively coupled over a network 804 to the other  
9 servers 806 and one or more databases 808. One or more of the other servers 806  
10 are also respective domain controllers for managing user access to directory  
11 resources.

12 In one implementation, the provisioning server 802 is operatively coupled  
13 to the network through one or more server appliances (not shown) located on the  
14 extreme outside of a server farm 810 such as a Web server farm, a corporate  
15 portal, or the like.

16 The provisioning server 802 includes a processor coupled to a memory.  
17 (Exemplary aspects of the processor and memory are described in greater detail in  
18 reference to processing unit 1132 and system memory 1134 in reference to Fig.  
19 11). The processor is configured to fetch execute computer-executable  
20 instructions from application programs 814 and configured to fetch data from  
21 program data 822. The application programs include a directory service 816, a  
22 provisioning service 818, and other applications 820 such as an operating system  
23 and text string editing application (e.g., an XML editor).

24 The directory service 816 stores information about objects such as  
25 computers and resources on a network 804 and makes this information available to

1 users and network administrators in a directory that ties disparate databases 808,  
2 or “directories” of data together into a single, logical directory, or “metadirectory”.  
3 The only objects that can be represented in the distributed directory are those that  
4 meet the content class qualifications specified by the directory schema 600. Such  
5 databases 808 or directories include, for example, directories of enterprise users,  
6 resources, financial information, corporate e-mail systems, network operating  
7 systems, and the like. A database is an object-oriented database such as an XML  
8 database, a Hypertext Markup Language (HTML) database, an SQL server  
9 database, and so on.

10 The provisioning service 818 is an instance object of the provisioning  
11 service 718 class of Fig. 7. The provisioning service interacts with the directory  
12 service 816 to coordinate and manage workflows based on an integrated view of  
13 directory resource identity. One aspect of the directory service/provisioning  
14 service interaction is the respective propagation (i.e., the directory service) and  
15 corresponding receipt (i.e., the provisioning service) of directory-based events that  
16 correspond to directory object status changes. (An exemplary procedure 900  
17 performed by the provisioning service is described in greater detail below in  
18 reference to Figs. 9 and 10).

19 The provisioning service instantiates one or more workflow objects (based  
20 on content class 710), provisioning status objects (based on content class 712),  
21 service history objects (based on content class 714), provisioning service class  
22 objects (based on content class 718) based on a state change to an object in the  
23 directory. To monitor or detect directory-based events, the provisioning service  
24 818 of Fig. 8 can be a Win32 service that monitors Lightweight Data Access  
25 Protocol (LDAP) directory synchronization (“Dirsync”) control associated with an

1 Active Directory® naming context to detect directory based events. (Win32 is a  
2 well-known WINDOWS® application programming interface (API). Techniques  
3 to interface with the LDAP directory synchronization control are well-known in  
4 the art of computer programming.)

5 Program data 822 includes the provisioning enabled directory schema 600,  
6 one or more workflow documents 826, a workflow schema 828, one or more event  
7 association documents 830, an event association schema 832, and one or more  
8 status monitoring documents 834 (for recording workflow status).

9 The particular workflows 826 to execute is/are determined by mapping the  
10 status change to one or more workflows using event association objects (based on  
11 content class 720) stored in the directory. Respective logic of the event  
12 association objects are represented by event association documents 830, the  
13 semantics of which are regulated by the event association schema 832. An event  
14 evaluation object indicates a set of declarative conditions that categorize/map a  
15 detected event to a defined workflow object 826. The declarative conditions are  
16 based on a type of event (e.g., an add event, a delete event, a rename event, and  
17 the like), a scope of the event (what portion of the directory tree the event occurred  
18 in), any filtering criteria (such as the structural content classes 710 through 720 of  
19 Fig. 7 or attribute values that are pertinent to the event), and/or the like.

20 An administrator defines the event/workflow mapping 830 as a text string  
21 such as an XML string. (See, for example, the XML event evaluation strings of  
22 Tables 2 and 3). Such a text string is “assigned to” a particular event evaluation  
23 object instance, which in turn is represented in the directory as an object/resource.  
24 Any number of various event evaluation objects can be defined and persisted in  
25 the directory. The semantics of event evaluation (e.g., such as the semantics

1 expressed in Tables 2 and 3) are based on the event evaluation schema 832, which  
2 defines a number of allowable declarative conditions and/or sequences that can be  
3 used to determine if an event maps to a particular workflow. Thus, the event  
4 evaluation schema is used as part of event evaluation text string validation that  
5 determines if a defined event evaluation includes pre-defined elements/tags.

6 The sequences of tasks and information corresponding to the mapped  
7 workflows are represented by workflow documents 826. A workflow object  
8 indicates a set of declarative conditions and/or sequence of actions that correspond  
9 to the identified workflow. An administrator defines the conditions and/or  
10 sequences of actions that correspond to a workflow. As discussed, this workflow  
11 definition is a text string such as an XML string. (See, for example, the XML  
12 workflow string of Table 4). This text string is “assigned to” a particular  
13 workflow object instance, which is a directory object/resource. Any number of  
14 various workflow definitions can be defined and persisted in the directory.

15 The semantics of a workflow are based on the workflow schema 828,  
16 which defines a number of allowable declarative conditions and/or sequences that  
17 can be defined in a workflow string that is put on a corresponding workflow object  
18 826. Thus, the workflow schema is used as part of workflow validation that  
19 determines if a defined workflow includes pre-defined elements/tags.

20 Once a particular workflow has been mapped to the directory object’s status  
21 change, an object instance of the workflow base content class 710 can be  
22 generated with the corresponding sequence of tasks and data place onto the  
23 object’s corresponding flexible attribute(s).

24 A provisioning status monitoring object 834 is an instance object of the  
25 status monitoring structural class 712 of Fig. 7. An instance of the provisioning

1 status object 712 is generated and used by the provisioning service/agent 818 to  
2 monitor a workflow's 826 status, and to reflect the current and past conditions  
3 within a directory that are attributed to the workflow 826 (e.g., as any directory  
4 resources affected by the workflow). The properties of this class indicate the  
5 sequences of actions that makeup a workflow along with status information that  
6 corresponds to the sequences.

7 The generation of provisioning status objects, service history objects, and  
8 so on, is policy based.

9 Convergence of Workflow to a Desired Directory State

10 The system 800 provides a mechanism to converge the results of a  
11 sequence of tasks that define a workflow to obtain a desired directory state.  
12 Specifically, the provisioning agent 818 identifies each workflow that executes and  
13 the identity of the directory resource(s) that have been affected by a workflow.  
14 Additionally, any conditional aspects of a workflow that may depend on any un-  
15 implemented aspects of the workflow can be satisfied because if a computer or  
16 other resource used in workflow execution is not available (e.g., malfunctions, is  
17 taken off-line, etc.), the provisioning agent 818 records the unavailability and non-  
18 occurrence of the corresponding workflow operation(s) in the associated workflow  
19 status monitoring object 834. The provisioning service 818 can continuously retry  
20 failed or otherwise incomplete tasks of a workflow. If an unavailable resource  
21 becomes again available (e.g., back on-line) after its unavailability, the appropriate  
22 declarative conditions and operational sequences defined on the workflow object  
23 can ensure that the previously failed operation completes.

1 In light of this, the described provisioning system 800 provides a  
2 substantial level of self-correcting behavior while implementing organizational  
3 policies. This is especially important to consider in a distributed computing  
4 environment, wherein more than one computer, and possibly dozens, hundreds, or  
5 even thousands of computers may be involved in executing any one operation of  
6 the workflow.

7 **Exemplary Procedure to Coordinate and Manage Provisioning**

8 Fig. 9 is a block diagram that shows an exemplary procedure to manage  
9 provisioning based on an integrated view of resource identity. At block 910, the  
10 procedure defines a number of workflows. For example, an administrator may  
11 define one or more workflow text strings, which are put onto respective workflow  
12 objects 826 as described above. (See, also workflow class 710 of Fig. 7, and  
13 Table 4). At block 912, the procedure generates a number of event association  
14 objects 828, which are stored or persisted in the directory. (See, also event  
15 association class 720 of Fig. 7, and Tables 2 and 3).

16 At block 914, the procedure, receives an event that corresponds to a status  
17 change of a directory object. Specifically, the provisioning service 818 interacts  
18 with the directory service 816 to monitor events that occur in the directory such as  
19 events that occur in any one of a number of different naming contexts (e.g., a  
20 DNC, an ANC, and/or the like) in a metadirectory. Such events include directory-  
21 based events such as the addition, modification, deletion, or renaming of a  
22 directory object. The provisioning service also detects events other than directory-  
23 based events such as time-based events, user actions (e.g., clicking a mouse button  
24 or pressing a key), system occurrences (e.g., running out of memory), and the like.

1 Accordingly, the types of events monitored by the provisioning service are  
2 customizable and extensible.

3 Responsive to receiving the event, at block 916, the procedure evaluates the  
4 event to map it to a workflow. To accomplish this, the provisioning service (or  
5 "provisioning agent") 818 categorizes the event to determine if there is an  
6 associated workflow that needs to be executed in response to the event. At block  
7 918, the procedure determines if the event corresponds to a defined workflow  
8 based on the evaluation of block 916. If so, at block 920, the procedure  
9 determines if an instance of the workflow object based on the scope of the event  
10 (i.e., those directory resources affected or that will be affected by the event) is  
11 already implemented. If an instance of the workflow based on the scope of the  
12 event has already been instantiated (block 920), the procedure continues at block  
13 810.

14 Otherwise, at block 922, the procedure 900 having determined that an  
15 instance of the workflow 826 based on the scope of the event has not yet been  
16 instantiated (block 920), generates an instance of the workflow object and places  
17 the identified workflow 826 onto the flexible attribute of the object. And, at block  
18 924, the procedure optionally generates a corresponding provisioning status  
19 monitoring object 834 of Fig. 8 to monitor the status of the newly instantiated  
20 workflow (block 922). The procedure continues at block 1010 of Fig. 9.

21 Fig. 10 shows further aspects of an exemplary procedure to manage  
22 provisioning based on an integrated view of resource identity. At block 1010, the  
23 procedure updates the identified workflow's (see, blocks 920 and 922 of Fig. 9)  
24 status monitoring object (block 924 of Fig. 9) to indicate that the particular event  
25 was received. Such a status update may involve evaluating the sequence of

operations that define the workflow (i.e., those stored on the monitoring object) to indicate workflow progress. For example, receipt of the event (block 914 of Fig. 9) may indicate the end/beginning of a particular workflow operation, the satisfaction of a declarative condition identified in the workflow, and/or the like. This indication is stored on the status object.

At block 1012, the procedure receives a second event. This second event corresponds to the provisioning agent's update of the provisioning status monitoring object (block 1010). At block 1014, the procedure re-evaluates sequence of operations and/or declarative conditions that make-up the workflow to determine the next step of the workflow to implement. In this manner, the provisioning service dynamically executes a workflow in response to directory events and coordinates and manages the provisioning process.

### Exemplary Computing Environment

Fig. 11 illustrates an example of a suitable operating environment 1120 in which an exemplary provisioning service may be implemented. Specifically, the exemplary provisioning service(s) described herein is implemented (wholly or in part) by any program module 1160-462 and/or operating system 1158 or a portion thereof. Exemplary computing environment 1120 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of an exemplary provisioning service. Neither should the computing environment 1120 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 1120.

1       The exemplary provisioning service is operational with numerous other  
2 general purpose or special purpose computing system environments or  
3 configurations. Examples of well known computing systems, environments,  
4 and/or configurations that may be suitable for use with an exemplary provisioning  
5 service include, but are not limited to, personal computers, server computers, thin  
6 clients, thick clients, hand-held or laptop devices, multiprocessor systems,  
7 microprocessor-based systems, set top boxes, programmable consumer electronics,  
8 network PCs, minicomputers, mainframe computers, distributed computing  
9 environments such as server farms and corporate intranets, and the like, that  
10 include any of the above systems or devices.

11      An exemplary provisioning service may be described in the general context  
12 of computer-executable instructions, such as program modules, being executed by  
13 a computer. Generally, program modules include routines, programs, objects,  
14 components, data structures, etc., that performs particular tasks or implement  
15 particular abstract data types. An exemplary provisioning service may also be  
16 practiced in distributed computing environments where tasks are performed by  
17 remote processing devices that are linked through a communications network. In  
18 a distributed computing environment, program modules may be located in both  
19 local and remote computer storage media including memory storage devices.

20      As shown in Fig. 11, the computing environment 1120 includes a general-  
21 purpose computing device in the form of a computer 1130. Computer 1130 is a  
22 suitable device for the implementation of the provisioning computer 802 of Fig. 8.  
23 The components of computer 1120 may include, by are not limited to, one or more  
24 processors or processing units 1132, a system memory 1134, and a bus 1136 that

1 couples various system components including the system memory 1134 to the  
2 processor 1132.

3 Bus 1136 represents one or more of any of several types of bus structures,  
4 including a memory bus or memory controller, a peripheral bus, an accelerated  
5 graphics port, and a processor or local bus using any of a variety of bus  
6 architectures. By way of example, and not limitation, such architectures include  
7 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)  
8 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)  
9 local bus, and Peripheral Component Interconnects (PCI) bus also known as  
10 Mezzanine bus.

11 Computer 1130 typically includes a variety of computer readable media.  
12 Such media may be any available media that is accessible by computer 1130, and  
13 it includes both volatile and non-volatile media, removable and non-removable  
14 media. In Fig. 11, the system memory includes computer readable media in the  
15 form of volatile memory, such as random access memory (RAM) 1140, and/or  
16 non-volatile memory, such as read only memory (ROM) 1138. A basic  
17 input/output system (BIOS) 1142, containing the basic routines that help to  
18 transfer information between elements within computer 1130, such as during start-  
19 up, is stored in ROM 1138. RAM 1140 typically contains data and/or program  
20 modules that are immediately accessible to and/or presently be operated on by  
21 processor 1132.

22 Computer 1130 may further include other removable/non-removable,  
23 volatile/non-volatile computer storage media. By way of example only, Fig. 11  
24 illustrates a hard disk drive 1144 for reading from and writing to a non-removable,  
25 non-volatile magnetic media (not shown and typically called a “hard drive”), a

1 magnetic disk drive 1146 for reading from and writing to a removable, non-  
2 volatile magnetic disk 1148 (e.g., a “floppy disk”), and an optical disk drive 1150  
3 for reading from or writing to a removable, non-volatile optical disk 1152 such as  
4 a CD-ROM, DVD-ROM or other optical media. The hard disk drive 1144,  
5 magnetic disk drive 1146, and optical disk drive 1150 are each connected to bus  
6 1136 by one or more interfaces 1154.

7 The drives and their associated computer-readable media provide  
8 nonvolatile storage of computer readable instructions (e.g., the application  
9 programs 814 of Fig. 8), data structures, program modules, and other data (e.g.,  
10 the program data 822 of Fig. 8) for computer 1130. Although the exemplary  
11 environment described herein employs a hard disk, a removable magnetic disk  
12 1148 and a removable optical disk 1152, it should be appreciated by those skilled  
13 in the art that other types of computer readable media which can store data that is  
14 accessible by a computer, such as magnetic cassettes, flash memory cards, digital  
15 video disks, random access memories (RAMs), read only memories (ROM), and  
16 the like, may also be used in the exemplary operating environment.

17 A number of program modules may be stored on the hard disk, magnetic  
18 disk 1148, optical disk 1152, ROM 1138, or RAM 1140, including, by way of  
19 example, and not limitation, an operating system 1158, one or more application  
20 programs 1160, other program modules 1162 (e.g., a directory service 816 and a  
21 provisioning service agent 818), and program data 1164 (e.g., a provisioning  
22 enabled directory schema 824, a workflow document 826, a workflow schema  
23 828, an event association document 830, an event association schema 832, and a  
24 status monitoring document 834).

1        Each of such operating system 1158, one or more application programs  
2        1160, other program modules 1162, and program data 1164 (or some combination  
3        thereof) may include an embodiment of an exemplary provisioning service. More  
4        specifically, each may include an embodiment of a provisioning service 818, and a  
5        directory schema 600 that includes a plurality of provisioning objects for  
6        managing a workflow instantiated by the provisioning service. The provisioning  
7        service includes a number of services such as an event association service, a  
8        workflow process, a workflow status-monitoring process, and the like. The  
9        provisioning objects in the directory schema include event association objects,  
10      workflow objects, status-monitoring objects, and the like.

11      A user may enter commands and information into computer 1130 through  
12      input devices such as keyboard 1166 and pointing device 1168 (such as a  
13      “mouse”). Other input devices (not shown) may include a microphone, joystick,  
14      game pad, satellite dish, serial port, scanner, or the like. These and other input  
15      devices are connected to the processing unit 1132 through a user input interface  
16      1170 that is coupled to bus 1136, but may be connected by other interface and bus  
17      structures, such as a parallel port, game port, or a universal serial bus (USB).

18      A monitor 1172 or other type of display device is also connected to bus  
19      1136 via an interface, such as a video adapter 1174. In addition to the monitor,  
20      personal computers typically include other peripheral output devices (not shown),  
21      such as speakers and printers, which may be connected through output peripheral  
22      interface 1175.

23      Computer 1130 may operate in a networked environment using logical  
24      connections to one or more remote computers, such as a remote computer 1182.  
25      Remote computer 1182 may include many or all of the elements and features

described herein relative to computer 1130. Logical connections include a local area network (LAN) 1177 and a general wide area network (WAN) 1179. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

When used in a LAN networking environment, the computer 1130 is connected to LAN 1177 via network interface or adapter 1186. When used in a WAN networking environment, the computer typically includes a modem 1178 or other means for establishing communications over the WAN 1179. The modem 1178, which may be internal or external, may be connected to the system bus 1136 via the user input interface 1170 or other appropriate mechanism.

Depicted in Fig. 11, is a specific implementation of a WAN via the Internet. Computer 1130 typically includes a modem 1178 or other means for establishing communications over the Internet 1180. Modem 1178, which may be internal or external, is connected to bus 1136 via interface 1170.

In a networked environment, program modules depicted relative to the personal computer 1130, or portions thereof, may be stored in a remote memory storage device. By way of example, and not limitation, Fig. 11 illustrates remote application programs 1189 as residing on a memory device of remote computer 1182. It will be appreciated that the network connections shown and described are exemplary and other means of establishing a communications link between the computers may be used.

## **Computer-Executable Instructions**

An implementation of an exemplary provisioning service may be stored on or transmitted across some form of computer readable media. Computer readable

1 media can be any available media that can be accessed by a computer. By way of  
2 example, and not limitation, computer readable media may comprise “computer  
3 storage media” and “communications media.”

4 “Computer storage media” include volatile and non-volatile, removable and  
5 non-removable media implemented in any method or technology for storage of  
6 information such as computer readable instructions, data structures, program  
7 modules, or other data. Computer storage media includes, but is not limited to,  
8 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,  
9 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic  
10 tape, magnetic disk storage or other magnetic storage devices, or any other  
11 medium which can be used to store the desired information and which can be  
12 accessed by a computer.

13 “Communication media” typically embodies computer readable  
14 instructions, data structures, program modules, or other data in a modulated data  
15 signal, such as carrier wave or other transport mechanism. Communication media  
16 also includes any information delivery media.

17 The term “modulated data signal” means a signal that has one or more of its  
18 characteristics set or changed in such a manner as to encode information in the  
19 signal. By way of example, and not limitation, communication media includes  
20 wired media such as a wired network or direct-wired connection, and wireless  
21 media such as acoustic, RF, infrared, and other wireless media. Combinations of  
22 any of the above are also included within the scope of computer readable media.

1      **Conclusion**

2      Although the provisioning service based on directory technology has been  
3      described in language specific to structural features and/or methodological steps, it  
4      is to be understood that the provisioning service based on directory technology  
5      defined in the appended claims is not necessarily limited to the specific features or  
6      steps described. Rather, the specific features and steps are disclosed as preferred  
7      forms of implementing the claimed present invention.

8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25